arXiv:1511.05803v1 [math.NA] 18 Nov 2015

Tractability of Multivariate Problems for Standard and Linear Information in the Worst Case Setting: Part I

Erich Novak*
Mathematisches Institut, Universität Jena
Ernst-Abbe-Platz 2, 07743 Jena, Germany
email: erich.novak@uni-jena.de
Henryk Woźniakowski†
Department of Computer Science, Columbia University,
New York, NY 10027, USA, and
Institute of Applied Mathematics, University of Warsaw
ul. Banacha 2, 02-097 Warszawa, Poland
email: henryk@cs.columbia.edu

November 19, 2015

Abstract

We present a lower error bound for approximating linear multivariate operators defined over Hilbert spaces in terms of the error bounds for appropriately constructed linear functionals as long as algorithms use function values. Furthermore, some of these linear functionals have the same norm as the linear operators. We then apply this error bound for linear (unweighted) tensor products. In this way we use negative tractability results known for linear functionals to conclude the same negative results for linear operators. In particular, we prove that L_2-multivariate approximation defined for standard Sobolev space suffers the curse of dimensionality if function values are used although the curse is not present if linear functionals are allowed.

*This author was partially supported by the DFG-Priority Program 1324.

†This author was partially supported by the National Science Foundation and by the National Science Centre, Poland, based on the decision DEC-2013/09/B/ST1/04275.

1 Introduction

The understanding of the intrinsic difficulty of approximation of d-variate problems is a challenging problem especially when d is large. We consider algorithms that approximate d-variate problems and use finitely many linear functionals: we compare the class $\Lambda^{\rm all}$ of arbitrary linear information functionals with the class $\Lambda^{\rm std}$ of information functionals that are given by function evaluations at single points.

To find best algorithms for the class $\Lambda^{\rm all}$ is usually much easier than for the class $\Lambda^{\rm std}$, in particular if the source space is a Hilbert space. This is especially the case for the worst case setting. The state of art may be found in [9], where the reader may find a number of surprising results. For example, there are multivariate problems for which the best rate of convergence of algorithms using n appropriately chosen linear functionals is $n^{-1/2}$ whereas for n function values the best rate can be arbitrarily bad, i.e., like $1/\ln(\ln(\cdots\ln(n))))$, where the number of ln can be arbitrarily large, see [4] which is also reported in [9] pp. 292-304. Furthermore, the dependence on d may be quite different for the linear and standard classes. There are examples of interesting multivariate problems for which the dependence on d *is not* exponential for the class $\Lambda^{\rm all}$, and *is* exponential for the class $\Lambda^{\rm std}$. The exponential dependence on d is called the *curse of dimensionality*. On the other hand, for some other multivariate problems there is no difference between $\Lambda^{\rm all}$ and $\Lambda^{\rm std}$. Examples can be found, in particular, in [2] and [6, 7, 9].

Tractability deals with how the intrinsic difficulty of a multivariate problem depends on d and on ε^{-1}, where ε is an error threshold. We would like to know when the curse of dimensionality holds and when we have a specific dependence on d which is not exponential. There are various ways of measuring the lack of exponential dependence and that leads to different notions of tractability. In particular, we have polynomial tractability (PT) if the intrinsic difficulty is polynomial in both d and ε^{-1}. We have quasi-polynomial tractability (QPT) if the intrinsic difficulty is at most proportional to $\varepsilon^{-t\,\ln d}$ for some t independent of d and ε.

Obviously, tractability may depend on which of the classes $\Lambda^{\rm std}$ or $\Lambda^{\rm all}$ is used. Tractability results for $\Lambda^{\rm std}$ cannot be better than for $\Lambda^{\rm all}$. The main question is for which multivariate problems they are more or less the same or for which multivariate problems they are essentially different.

These questions were already addressed in [6, 7, 9]. Still, especially the worst case setting is not fully understood. We would like to get a better understanding how the power of the standard class $\Lambda^{\rm std}$ is related to the power of the class $\Lambda^{\rm all}$ of information. Ideally, we would like to characterize for which multivariate problems the classes $\Lambda^{\rm std}$ and $\Lambda^{\rm all}$ lead to more or less the same tractability results and for which tractability results are essentially different.

We plan to write a number of papers about this problem under the same title. We present

the first part of this project. We restrict ourselves to linear multivariate problems defined as approximation of a linear continuous operator $S : F \to G$ for general Hilbert spaces F and G. Since we want to study the class $\Lambda^{\rm std}$ we need to assume that function values are well defined and they correspond to linear continuous functionals. This is equivalent to assuming that F is a reproducing kernel Hilbert space.

For the worst case setting and for the class $\Lambda^{\rm all}$, it is known what is the best way to approximate S. The intrinsic difficulty of approximating S is defined as the *information complexity* which is the minimal number of linear functionals which are needed to find an algorithm whose worst case error is at most $\varepsilon\|S\|$. This depends on the eigenvalues of the operator $S^*S : F \to F$. For the class $\Lambda^{\rm std}$ the situation is much more complex and the *information complexity*, which is now the minimal number of function values needed to get an error $\varepsilon\|S\|$, depends not only on the eigenvalues of S^*S.

Our first result is the construction of continuous linear functionals I which are at most as hard to approximate as S for the class $\Lambda^{\rm std}$. Furthermore, we characterize I for which $\|I\| = \|S\|$. They are of the form

$$I = \langle \cdot, S^* g \rangle_F \quad \text{with} \quad g = \lambda_1^{-1/2} S\eta,$$

where λ_1 is the largest eigenvalue of S^*S and η of norm 1 belongs to the eigenspace corresponding to λ_1. Hence, if λ_1 is of multiplicity 1 then the choice of g is essentially unique. If λ_1 is of multiplicity larger than 1, then the choice of g is not unique and may lead to trivial or hard linear functionals I.

For I with $\|I\| = \|S\|$, the information complexity of I for the class $\Lambda^{\rm std}$ is at most equal to the information complexity of S. Hence, if I is hard to approximate so is S.

The essence of this result is that for approximation of linear functionals over some Hilbert spaces there is a proof technique which allows to find sharp error bounds. This proof technique was developed in [5] and requires that the reproducing kernel of F_1 has a so called decomposable part.

We verify how this lower bound on approximating S works for linear d-folded (unweighted) tensor product problems. Then the corresponding linear functionals I are also d-folded tensor products. We then may apply the existing negative tractability results for I and conclude the same negative tractability results for S.

We illustrate our approach for a number of examples. In particular, we consider the Sobolev space $F_d = F_1^{\otimes d}$ with the reproducing kernel

$$K_d(x,t) = \prod_{j=1}^{d} (1 + \min(x_j, t_j)) \quad \text{for all} \quad x, t \in [0,1]^d,$$

and $G_d = L_2([0,1]^d)$.

Let $S = S_d : F_d \to G_d$ be any non-zero tensor product operator $S_d = S_1^{\otimes d}$ with $S_1 : F_1 \to G_1$. Let $\{\lambda_j\}$ be the ordered sequence of eigenvalues of $S_1^* S_1$. Let

$$\text{decay}_\lambda := \sup\{\, r > 0 \,:\, \lim_{n\to\infty} n^r \lambda_n = 0 \,\}$$

denote the polynomial decay of the eigenvalues λ_n. If the set of r above is empty we set $\text{decay}_\lambda = 0$.

Let $\mathbb{S} = \{S_d\}_{d=1}^\infty$. It is known, see [2], that $\mathbb{S}$ is quasi-polynomially tractable (QPT) for the class Λ^{all} iff $\lambda_2 < \lambda_1$ and $\text{decay}_\lambda > 0$. Furthermore, if λ_2 is positive then $\mathbb{S}$ is not polynomially tractable (PT). On the other hand, if $\lambda_2 = \lambda_1$ then $\mathbb{S}$ suffers from the curse of dimensionality for the class Λ^{all} (and obviously also for Λ^{std}).

For the class Λ^{std}, assume without loss of generality that $\lambda_2 < \lambda_1$. Let η_1 be a normalized eigenfunction corresponding to the largest eigenvalue λ_1. We prove that $\mathbb{S}$ suffers from the curse of dimensionality if

$$\eta_1 \neq aK_1(\cdot, t) = a\,(1 + \min(\cdot, t)) \quad \text{for all} \;\; a \in \mathbb{R} \;\; \text{and} \;\; t \in [0,1]. \tag{1}$$

This means that we have the curse of dimensionality as long as the eigenfunction of S_1 corresponding to the largest eigenvalue is not proportional to the univariate reproducing kernel with one argument fixed. We then verify that this assumption holds for multivariate approximation, i.e., for $S_1 f = f$. This partially solves the open problem 131 from [9] p. 361.

We believe that the assumption (1) is also necessary for the curse. More generally, we believe that for $\eta_1 = aK_1(\cdot, t)$ for some real a and t from the common domain of univariate functions, and of course for $\lambda_2 < \lambda_1$ and $\text{decay}_\lambda > 0$, we have QPT for the class Λ^{std} and this holds for any K_1. But this will be the subject of the next part of our project.

In this paper we discuss only unweighted tensor products and that is why we do not have polynomial tractability (PT) for problems with two positive eigenvalues. PT and other notions of tractability may hold if we consider weighted tensor products with sufficiently decaying weights. This will be also a subject of our next study.

2 Relation between Linear Functionals and Operators

Consider a continuous linear and non-zero operator $S : F \to G$, where F is a reproducing kernel Hilbert space of real functions f defined over a common domain $D \subset \mathbb{R}^d$ for some positive integer d, and G is a Hilbert space. We approximate S by algorithms A_n that use at most n linear functionals. Without loss of generality we may assume that A_n is linear, see e.g., [6, 12]. That is,

$$A_n f = \sum_{j=1}^n L_j(f)\, g_j$$

for some $L_j \in F^*$ and $g_j \in G$. Using the same proof as in [6] p. 345, we may also assume that $g_j = Sf_j$ for some $f_j \in F$.

We consider two classes of linear functionals L_j's:

- the linear class of information $\Lambda^{\rm all}$ which consists of all continuous and linear functionals L_j's, i.e., $\Lambda^{\rm all} = F^*$, and
- the standard class of information $\Lambda^{\rm std}$ which consists of function values, i.e., $L_j(f) = \langle f, K(\cdot, t_j)\rangle_F = f(t_j)$ for some $t_j \in D$, where $K : D \times D \to \mathbb{R}$ is the reproducing kernel of F.

The nth minimal (worst case) error of approximating S for the class $\Lambda \in \{\Lambda^{\rm std}, \Lambda^{\rm all}\}$ is defined as

$$e_n(S, \Lambda) = \inf_{L_1,\ldots,L_n \in \Lambda,\, g_1,\ldots,g_n \in G} \sup_{\|f\|_F \le 1} \|Sf - A_n f\|_G.$$

For $n = 0$, we take $A_n = 0$ and then we obtain the initial error which is independent of Λ and given by

$$e_0(S, \Lambda) = e_0(S) = \|S\|_{F\to G}.$$

For the class $\Lambda^{\rm all}$, it is well known that $\lim_{n\to\infty} e_n(S, \Lambda^{\rm all}) = 0$ iff S is compact. This is why we always may assume that S is compact. Then it is known that the nth minimal errors depend on the eigenvalues of

$$W = S^*S : F \to F.$$

More precisely, let (λ_j, η_j) be eigenpairs of W,

$$W\eta_j = \lambda_j \eta_j, \quad \text{with} \quad \langle \eta_j, \eta_k\rangle_F = \delta_{j,k} \ \text{ and } \ \lambda_1 \ge \lambda_2 \ge \cdots .$$

Observe here that the λ_j are uniquely defined, but the η_j are not unique. Moreover, we formally define $\lambda_j = 0$ if the dimension of F is finite and j is larger than this dimension. Then

$$e_n(S, \Lambda^{\rm all}) = \sqrt{\lambda_{n+1}} \quad \text{for } \ n = 0, 1, \ldots .$$

Hence $e_0(S) = \|S\|_{F\to G} = \sqrt{\lambda_1}$, and since S is non-zero we have $\lambda_1 > 0$.

The situation is much more complicated for the class $\Lambda^{\rm std}$. Obviously,

$$e_n(S, \Lambda^{\rm std}) \ge e_n(S, \Lambda^{\rm all}) \quad \text{for all } \ n = 0, 1, \ldots,$$

but it is not clear when the sequences $e_n(S, \Lambda^{\rm std})$ and $e_n(S, \Lambda^{\rm all})$ behave similarly. There are many papers studying the powers of $\Lambda^{\rm std}$ and $\Lambda^{\rm all}$. The state of art can be found in [9].

In this paper we continue this study and show that the sequence $e_n(S, \Lambda^{\rm std})$ can behave quite differently than the sequence $e_n(S, \Lambda^{\rm all})$. This will be done by showing first that many continuous and linear functionals are at most as hard to approximate as S and there are functionals for which we can also match the initial error $e_0(S)$ of S.

More precisely, for any $g \in G$ with $\|g\|_G = 1$ define

$$I_g f = \langle f, S^* g \rangle_F \qquad \text{for all} \quad f \in F.$$

Note that $S^* g \in F$ and therefore I_g is a continuous linear functional. The nth minimal error of approximating I_g is defined as for S, this time with G replaced by $\mathbb{R}$. Clearly, $e_0(I_g) = \|S^* g\|_F$.

Theorem 1.

For any $g \in G$ with $\|g\|_G = 1$ we have

$$e_n(I_g, \Lambda^{\rm std}) \le e_n(S, \Lambda^{\rm std}) \qquad \text{for all} \quad n \in \mathbb{N}_0.$$

Furthermore,

$$e_0(I_g) = e_0(S) \quad \text{and} \quad \|g\|_G = 1 \qquad \text{iff} \qquad g = \lambda_1^{-1/2} S\eta,$$

where η is any element of norm 1 with $S^* S\eta = \lambda_1 \eta$. □

Proof. Take an arbitrary linear algorithm $A_n f = \sum_{j=1}^n f(t_j)\, S f_j$ for approximating S. Define

$$B_n f = \sum_{j=1}^n f(t_j)\ \langle f_j, S^* g \rangle_F$$

as a linear algorithm for approximating I_g. Then

$$I_g f - B_n f = \left\langle f - \sum_{j=1}^n f(t_j)\, f_j, S^* g \right\rangle_F = \left\langle S f - \sum_{j=1}^n f(t_j)\, S f_j, g \right\rangle_G = \langle S f - A_n f, g \rangle_G .$$

Therefore

$$|I_g f - B_n f| \le \|S f - A_n f\|_G \, \|g\|_G = \|S f - A_n f\|_G.$$

Taking the supremum over the unit ball of F and then the infimum over t_j's and f_j's, we conclude that

$$e_n(I_g, \Lambda^{\rm std}) \le e_n(S, \Lambda^{\rm std}),$$

as claimed.

Let m be the multiplicity of the largest eigenvalue λ_1, i.e., $\mathrm{span}(\eta_1, \eta_2, \ldots, \eta_m)$ is the eigenspace of W for the eigenvalue λ_1. Take now $g = \lambda_1^{-1/2} S\eta$ for any $\eta \in \mathrm{span}(\eta_1, \eta_2, \ldots, \eta_m)$ with $\|\eta\|_F = 1$. Then we have $S^* g = \lambda_1^{-1/2} W\eta = \lambda_1^{1/2}\eta$ and

$$e_0(I_g) = \|S^* g\|_F = \sqrt{\lambda_1} = \|S\|_{F\to G} = e_0(S).$$

Furthermore,

$$\|g\|_G^2 = \left\langle \lambda_1^{-1/2} S\eta, \lambda_1^{-1/2} S\eta \right\rangle_G = \lambda_1^{-1} \langle \eta, W\eta \rangle_F = \lambda_1^{-1}\lambda_1 \|\eta\|_F^2 = 1.$$

We need to show that $e_0(I_g) = e_0(S)$ holds only for such g. Take then any g from G such that $\|g\|_G = 1$ and $\|S^* g\|_F = \lambda_1^{1/2}$. We can represent

$$g = \alpha\, S\eta + h,$$

where $\eta \in \mathrm{span}(\eta_1, \eta_2, \ldots, \eta_m)$ with $\|\eta\|_F = 1$, $\alpha \in \mathbb{R}$ and h is orthogonal to $S\eta_j$ for all $j = 1, 2, \ldots, m$, i.e.,

$$0 = \langle h, S\eta_j \rangle_G = \langle S^* h, \eta_j \rangle_F \qquad \text{for } \; j = 1, 2, \ldots, m.$$

Since $\|S\eta\|_G = \lambda_1^{1/2}$ we have

$$1 = \|g\|_G^2 = \alpha^2 \lambda_1 + \|h\|_G^2.$$

On the other hand,

$$\begin{aligned} 1 = \frac{1}{\lambda_1}\|S^* g\|_F^2 &= \frac{1}{\lambda_1}\|\alpha\, S^* S\eta + S^* h\|_F^2 = \frac{1}{\lambda_1}\|\alpha\lambda_1\, \eta + S^* h\|_F^2 \\ &= \frac{1}{\lambda_1}\left(\alpha^2\lambda_1^2 + \|S^* h\|_F^2\right) = \alpha^2\lambda_1 + \frac{1}{\lambda_1}\|S^* h\|_F^2. \end{aligned}$$

We now analyze $\|S^* h\|_F$. Note that $\|S^* h\|_F^2 = \langle h, SS^* h \rangle_G$. Let

$$G = \overline{SS^*(G)} \oplus [\overline{SS^*(G)}]^{\perp}.$$

Hence, for any h from G we have $h = h_1 + h_2$ with $h_1 \in \overline{SS^*(G)}$ and h_2 orthogonal to $\overline{SS^*(G)}$. Then

$$\|S^* h\|_F^2 = \langle h_1 + h_2, SS^* h_1 \rangle_G = \langle h_1, SS^* h_1 \rangle_G \,. \tag{2}$$

Let $k = \sup\{j : \ \lambda_j > 0\}$. Hence, if all $\lambda_j > 0$ then $k = \infty$, otherwise k is the number of positive eigenvalues λ_j. Clearly, $k \geq m$.

We know that $S^*S\eta_j = \lambda_j\eta_j$. Then $S\eta_j \neq 0$ for all finite j which are at most k. Multiplying the last equation by S we obtain

$$SS^*(S\eta_j) = \lambda_j(S\eta_j).$$

Hence, $(\lambda_j, \lambda_j^{-1/2}S\eta_j)$ is an eigenpair of SS^* and

$$\left\langle \lambda_j^{-1/2}S\eta_j, \lambda_i^{-1/2}S\eta_i \right\rangle_G = \frac{\langle S\eta_j, S\eta_i\rangle_G}{\sqrt{\lambda_j\lambda_i}} = \frac{\langle \eta_j, S^*S\eta_i\rangle_F}{\sqrt{\lambda_j\lambda_i}} = \frac{\lambda_i}{\sqrt{\lambda_j\lambda_i}}\langle \eta_j, \eta_i\rangle_F = \delta_{i,j}.$$

That means that $\lambda_j^{-1/2}S\eta_j$'s are orthonormal in G. Since $SS^*(G) \subset S(F)$, the $\lambda_j^{-1/2}S\eta_j$'s build a complete orthonormal system of $\overline{SS^*(G)}$ and, when we return to (2), we may write

$$h_1 = \sum_{j=1}^{k} \left\langle h_1, \lambda_j^{-1/2}S\eta_j \right\rangle_G \lambda_j^{-1/2}S\eta_j$$

and then

$$\langle h_1, SS^*h_1\rangle_G = \sum_{j=1}^{k} \left\langle h_1, \lambda_j^{-1/2}S\eta_j \right\rangle_G^2 \lambda_j.$$

Since $0 = \langle h, S\eta_j\rangle_G = \langle h_1, S\eta_j\rangle_G$ for all $j \leq m$, we conclude that

$$\begin{aligned} \langle h_1, SS^*h_1\rangle_G &= \sum_{j=m+1}^{k} \left\langle h_1, \lambda_j^{-1/2}S\eta_j \right\rangle_G^2 \lambda_j \leq \lambda_{m+1} \sum_{j=m+1}^{k} \left\langle h_1, \lambda_j^{-1/2}S\eta_j \right\rangle_G^2 \\ &= \lambda_{m+1}\|h_1\|_G^2 \leq \lambda_{m+1}\|h\|_G^2. \end{aligned}$$

From this, we get

$$\begin{aligned} 1 = \alpha^2\lambda_1 + \frac{1}{\lambda_1}\|S^*h\|_F^2 &\leq \alpha^2\lambda_1 + \frac{\lambda_{m+1}}{\lambda_1}\|h\|_G^2 = \alpha^2\lambda_1 + \|h\|_G^2 - \left(1 - \frac{\lambda_{m+1}}{\lambda_1}\right)\|h\|_G^2 \\ &= 1 - \left(1 - \frac{\lambda_{m+1}}{\lambda_1}\right)\|h\|_G^2. \end{aligned}$$

Since $\lambda_{m+1} < \lambda_1$ we conclude that $h = 0$ and $g = \alpha\, S\eta$ with $\alpha^2\lambda_1 = 1$. This completes the proof. □

For any g from G the linear functional I_g can be also written as

$$I_g f = \langle Sf, g\rangle_G \quad \text{for all} \quad f \in F.$$

Example 2. As an example, if we take $G = L_2([0,1]^d)$ then

$$I_g f = \int_{[0,1]^d} (Sf)(x)\, g(x)\, \mathrm{d}x.$$

Furthermore, if we additionally assume that F is continuously embedded in $L_2([0,1]^d)$ and take $S = \mathrm{APP}_d$ as multivariate approximation, $\mathrm{APP}_d f = f$ for all $f \in F$, then

$$I_g f = \int_{[0,1]^d} f(x)\, g(x)\, \mathrm{d}x.$$

If $g \equiv 1$ then

$$\mathrm{INT}_d f = I_g f = \int_{[0,1]^d} f(x)\, \mathrm{d}x$$

is multivariate integration, and

$$e_0(\mathrm{INT}_d) \le e_0(\mathrm{APP}_d).$$

This relation between multivariate integration and approximation has been used in many papers. For some spaces the norm of multivariate integration and approximation is the same. This is the case for Korobov spaces and some Sobolev spaces as will be reported later.

However, in general, the norm of multivariate integration is smaller and sometimes exponentially smaller than the norm of multivariate approximation. This is the case for some other Sobolev spaces. For instance, this holds for the space $F = F_d$ with the reproducing kernel

$$K_d(x,y) = \prod_{j=1}^{d} (1 + \min(x_j, y_j)) \quad \text{for all} \quad x_j, y_j \in [0,1].$$

It is known, see [9] pp. 353 and 411, that

$$e_0(\mathrm{INT}_d) = (4/3)^{d/2} = (1.3333\ldots)^{d/2} \quad \text{and} \quad e_0(\mathrm{APP}_d) = (1.35103388\ldots)^{d/2}.$$

Hence,

$$\frac{e_0(\mathrm{APP}_d)}{e_0(\mathrm{INT}_d)} = (1.013\ldots)^{d/2}.$$

Although $1.013\ldots$ is barely larger than one, the ratio of the initial errors for multivariate approximation and integration goes to infinity exponentially fast with d.

As we shall see, the multiplicity of the largest eigenvalue for this multivariate approximation is $m = 1$. Therefore, in order to match the norm of multivariate approximation we must

use a weighted integration problem $I_g(f) = (f, g)_{L_2}$ with $g = \lambda_1^{-1/2}\eta_1$ (or $g = -\lambda_1^{-1/2}\eta_1$) which for our example of F_d is not equal to the constant function 1. In Section 5 we will show that $g(x) = \prod_{j=1}^{d} g_1(x_j)$ for $x = [x_1, \ldots, x_d] \in [0,1]^d$. We find it interesting to know the "most difficult" integration problem I_g (with $\|g\|_2 = 1$) for a Hilbert space of functions and hence present the graph of the function g_1 in Figure 1. The same g is also the unique (up to a multiplicative constant) function that maximizes $\frac{\|g\|_2}{\|g\|_F}$ and hence solves an important optimization problem.

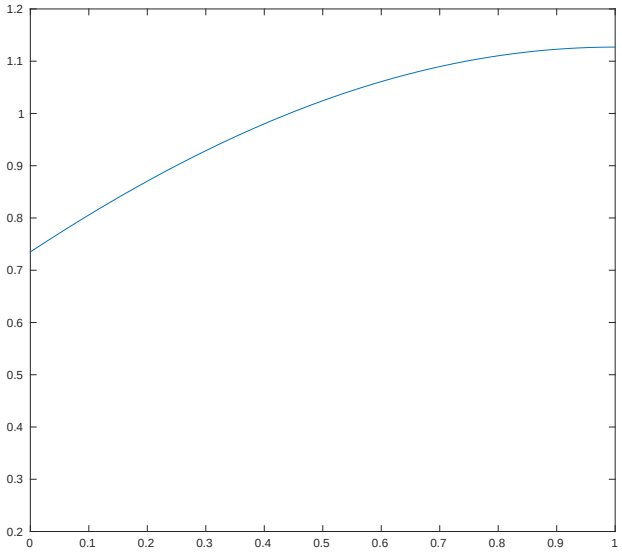

Figure 1: Density g for which $\|I_g\| = \|\mathrm{APP}_1\|$

We now show that the choice of η in $g = \lambda_1^{-1/2} S\eta$ may be important if the multiplicity of λ_1 is larger than 1. That is, it may happen that for some such g the functional I_g is trivial and for some other g, it may be very difficult.

Example 3.

Let F_1 be the space of functions $f : [0,1] \to \mathbb{R}$ that are constant over $[0, \frac{1}{2}]$ and $(\frac{1}{2}, 1]$. That is, for $f \in F_1$ there are real a and b such that $f(x) = a$ for all $x \in [0, \frac{1}{2}]$ and $f(x) = b$ for all $x \in (\frac{1}{2}, 1]$.

We equip F_1 with the L_2 norm which can be written (for the space F_1) as

$$\|f\|_{F_1} = \tfrac{1}{\sqrt{2}} \left(f^2(0) + f^2(1)\right)^{1/2}.$$

We define $F = G = F_1^{\otimes d}$ as the d folded tensor product of the space F_1. The space F consists of piecewise constant functions over 2^d subintervals of volume 2^{-d} which are a partition of the cube $[0,1]^d$. The space F_d is also equipped with the L_2 norm. Clearly, $\dim(F) = 2^d$.

Let $S : F \to G$ be the identity operator. Then $\lambda_j = 1$ for all $j = 1, 2, \ldots, 2^d$ and any nonzero function from F is an eigenfunction of $W = S^*S = I$. Clearly, $\lambda_{2^d+1} = 0$ and therefore

$$e_n(S, \Lambda^{\rm all}) = 1 \quad \text{for all} \;\; n < 2^d \;\; \text{and} \;\; e_{2^d}(S, \Lambda^{\rm all}) = 0.$$

Obviously, it also proves that $e_n(S, \Lambda^{\rm std}) = 1$ for $n < 2^d$ since $1 = e_n(S, \Lambda^{\rm all}) \le e_n(S, \Lambda^{\rm std}) \le e_0(S) = 1$.

For any g of norm 1 we have $e_0(I_g) = e_0(S) = 1$. We now show that $e_n(I_g, \Lambda^{\rm std})$ very much depends on the choice of g.

Suppose we take $g = 2^{d/2}$ over $[0, \frac{1}{2}]^d$ and $g = 0$ otherwise. Then

$$I_g(f) = 2^{d/2} \int_{[0,\frac{1}{2}]^d} f(x)\,{\rm d}x = 2^{-d/2}\, f(0)$$

is a trivial linear functional which can be solved exactly by using one function value at 0. Hence,

$$e_n(I_g, \Lambda^{\rm std}) = 0 \quad \text{for all} \;\; n \ge 1.$$

In this case, the bound $0 = e_n(I_g, \Lambda^{\rm std}) \le e_n(S, \Lambda^{\rm std}) = 1$ is useless.

Take now $g = 1$ over the cube $[0,1]^d$. Then

$$I_g(f) = \int_{[0,1]^d} f(x)\,{\rm d}x = \frac{1}{2^d} \sum_{j=[j_1,j_2,\ldots,j_d] \in \{0,1\}^d} f(t_{j_1}, t_{j_2}, \ldots, t_{j_d}),$$

where $t_0 = 0$ and $t_1 = 1$. We prove that the nth minimal error for $n \le 2^d$ is

$$e_n(I_d, \Lambda^{\rm std}) = \left(1 - n\,2^{-d}\right)^{1/2}.$$

Indeed, for $n = 2^d$ we can sample f at all points $f(t_1, t_2, \ldots, t_d)$ with $t_j \in \{0, 1\}$ and recover I_g exactly. Therefore $e_{2^d}(I_d, \Lambda^{\rm std}) = 0$.

Assume now that $n < 2^d$. Suppose we sample f at some $x_1, x_2, \ldots, x_n$ from the unit cube $[0,1]^d$. Then it is enough to take f which is zero at n sub-cubes that contain samples $x_1, x_2, \ldots, x_n$, and which takes a constant value c at $2^d - n$ sub-cubes. Taking c for which the norm is 1 we obtain the equation $1 = c^2(2^d - n)/2^d$ which yields $c = \pm\,(2^d/(2^d - n))^{1/2}$. Then

$$I_g(f) = c\,\frac{2^d - n}{2^d} = \pm\,\left(1 - n\,2^{-d}\right)^{1/2}.$$

All linear algorithms must approximate $I_g(f)$ by zero and therefore their worst case error is at least $\left(1 - n\,2^{-d}\right)^{1/2}$. The last bound is sharp if we take sample points $x_1, x_2, \ldots, x_n$ at disjoint sub-cubes, as claimed.

Hence, in this case we have

$$\left(1 - n\,2^{-d}\right)^{1/2} = e_n(I_g, \Lambda^{\rm std}) \le e_n(S, \Lambda^{\rm std}) = 1.$$

The bound is quite sharp as long as n is much smaller than 2^d.

□

For general spaces, we will use Theorem 1 for I_g with $g = \lambda_1^{-1/2} S\eta_1$. For the standard class of information $\Lambda^{\rm std}$, using lower bounds results for I_g from [7], we obtain lower bounds results for S. In this way we show, in particular, that we have sometimes the curse of dimensionality for $\Lambda^{\rm std}$ which is not present for the class $\Lambda^{\rm all}$.

3 Tractability Notions

We need to recall the definition of the information complexity for the so-called normalized error criterion. It is defined as the minimal number of linear functionals from the class Λ which are needed to reduce the initial error by a factor $\varepsilon \in (0,1)$, where $\Lambda \in \{\Lambda^{\rm std}, \Lambda^{\rm all}\}$. That is,

$$n(\varepsilon, S, \Lambda) = \min\{\, n \,:\, e_n(S, \Lambda) \le \varepsilon\, e_0(S)\}.$$

For the class $\Lambda^{\rm all}$, we obviously have

$$n(\varepsilon, S, \Lambda^{\rm all}) = \min\{\, n \,:\, \lambda_{n+1} \le \varepsilon^2\, \lambda_1\}.$$

Unfortunately, there is no such or similar formula for the class $\Lambda^{\rm std}$.

Assume now that we have a sequence

$$\mathbb{S} = \{S_d\}_{d=1}^{\infty}$$

of continuous linear non-zero operators $S_d : F_d \to G_d$, where F_d is a reproducing kernel Hilbert space of real function defined over $D_d \subset \mathbb{R}^d$ and G_d is a Hilbert space. In this case, we want to verify how the information complexity $n(\varepsilon, S_d, \Lambda)$ depends on ε^{-1} and d.

In this paper we will use only a few tractability notions which are defined as follows. We say that

- $\mathbb{S}$ suffers from the curse of dimensionality for the class $\Lambda \in \{\Lambda^{\rm std}, \Lambda^{\rm all}\}$ iff there are positive numbers c and C as well as $\varepsilon_0 \in (0,1)$ such that

$$n(\varepsilon, S_d, \Lambda) \geq C\,(1+c)^d$$

for all $\varepsilon \in (0, \varepsilon_0]$ and infinitely many d.

- $\mathbb{S}$ is quasi-polynomially tractable (QPT) for the class Λ iff there are non-negative numbers C and t such that

$$n(\varepsilon, S_d, \Lambda) \leq C\,\exp\big(t\,(1+\ln\,\varepsilon^{-1})(1+\ln\,d)\big) \qquad \text{for all } \ \varepsilon \in (0,1),\, d = 1,2,\ldots.$$

The infimum of numbers t satisfying the bound above is called the exponent of QPT and denoted by t^*.

- $\mathbb{S}$ is polynomially tractable (PT) for the class Λ iff there are nonnegative numbers C, p, q such that

$$n(\varepsilon, S_d, \Lambda) \leq C\,\varepsilon^{-p}\,d^q \qquad \text{for all } \ \varepsilon \in (0,1),\, d = 1,2,\ldots.$$

Clearly, PT implies QPT. More about these and other tractability concepts can be found in [6, 7, 9].

For the class $\Lambda^{\rm all}$, tractability notions depend on the decay of the eigenvalues $\lambda_{d,j}$ of the operator $W_d = S_d^* S_d : F_d \to F_d$. Necessary and sufficient conditions can be found in the works cited above. Again for the class $\Lambda^{\rm std}$, no such conditions are known and they cannot depend only on the eigenvalues $\lambda_{d,j}$.

4 Linear Tensor Product Problems

From now on we study a sequence

$$\mathbb{S} = \{S_d\}_{d=1}^{\infty}$$

of tensor product problems. Hence the spaces $F = F_d = F_1^{\otimes d}$ and $G = G_d = G_1^{\otimes d}$ as well as $S = S_d = S_1^{\otimes d}$ are given by tensor products of d copies of F_1 and G_1 as well as a continuous linear operator $S_1 : F_1 \to G_1$, respectively, where F_1 is a reproducing kernel Hilbert space of real univariate functions defined over $D_1 \subset \mathbb{R}$ and G_1 is a Hilbert space. Then F_d is a space of d-variate real functions defined on $D_d = D_1 \times D_1 \times \cdots \times D_1$ (d times).

An important example is given by multivariate approximation. That is, we now take $G_d = L_2([0,1]^d)$ and $S_1 = \mathrm{APP}_1 : F_1 \to G_1$ with the embedding operator $\mathrm{APP}_1 f = f$. Then

$S_d = \mathrm{APP}_d : F_d \to G_d$ is also the embedding operator $\mathrm{APP}_d f = f$ for all $f \in F_d$. In this case, we denote

$$\mathbb{S} = \mathbb{APP}, \quad \text{where} \quad \mathbb{APP} = \{\mathrm{APP}_d\}_{d=1}^{\infty}.$$

If K_1 is the reproducing kernel of F_1 then F_d is a reproducing kernel Hilbert space whose kernel is

$$K_d(x,y) = \prod_{j=1}^{d} K_1(x_j, y_j) \quad \text{for all} \;\; x = [x_1, x_2, \ldots, x_d],\; y = [y_1, y_2, \ldots, y_d] \in D_d.$$

For the class Λ^{all}, it is well known that the eigenpairs $(\lambda_{d,j}, \eta_{d,j})$ of $W_d = S_d^* S_d : F_d \to F_d$ are given in terms of the eigenpairs (λ_j, η_j) of the univariate operator $W_1 = S_1^* S_1 : F_1 \to F_1$. As before we assume that $\lambda_1 \geq \lambda_2 \geq \cdots$, $\langle \eta_j, \eta_k \rangle_{F_1} = \delta_{j,k}$ and that S_1 is non-zero. This means that $\lambda_1 > 0$. For the operator W_d we have

$$\{\lambda_{d,j}\}_{j=1}^{\infty} = \{\lambda_{j_1} \lambda_{j_2} \cdots \lambda_{j_d}\}_{j_1, j_2, \ldots, j_d = 1}^{\infty}.$$

Similarly, the eigenfunctions of W_d are of product form

$$\{\eta_{d,j}\}_{j=1}^{\infty} = \{\eta_{j_1} \otimes \eta_{j_2} \otimes \cdots \otimes \eta_{j_d}\}_{j_1, j_2, \ldots, j_d = 1}^{\infty},$$

where

$$[\eta_{j_1} \otimes \eta_{j_2} \otimes \cdots \otimes \eta_{j_d}](x) = \prod_{k=1}^{d} \eta_{j_k}(x_k) \quad \text{for all} \;\; x = [x_1, \ldots, x_d] \in D_d.$$

Then $\|S_d\|_{F_d \to G_d} = \|W_d\|_{F_d \to F_d}^{1/2} = \lambda_1^{d/2}$. Hence, the initial error is

$$e_0(S_d) = \lambda_1^{d/2}.$$

If $\lambda_2 = \lambda_1$ then we have at least 2^d eigenvalues of W_d equal to λ_1^d, and therefore

$$n(\varepsilon, S_d) \geq 2^d \quad \text{for all} \;\; \varepsilon \in (0,1) \;\; \text{and} \;\; d = 1, 2, \ldots .$$

In this case $\mathbb{S} = \{S_d\}_{d=1}^{\infty}$ suffers from the curse of dimensionality. On the other hand, it is proved in [2], see also [9] p. 112, that $\mathbb{S}$ is QPT for the class Λ^{all} iff $\lambda_2 < \lambda_1$ and

$$\mathrm{decay}_\lambda := \sup\{\, r > 0 \;:\; \lim_{n \to \infty} n^r \lambda_n = 0 \,\} > 0. \tag{3}$$

If the last conditions hold then the exponent of QPT is

$$t^* = \max\left(\frac{2}{\mathrm{decay}_\lambda}, \frac{2}{\ln \frac{\lambda_1}{\lambda_2}} \right).$$

Note that for $\lambda_2 = 0$ we have $t^* = 0$. In this case, S_d is a continuous linear functional and $n(\varepsilon, S_d, \Lambda^{\rm all}) = 1$ for all $\varepsilon \in [0,1)$ and all $d = 1, 2, \ldots$. If $\lambda_2 > 0$ then the exponent of QPT is positive and in this case it is also known that the problem $\mathbb{S}$ is *not* PT for the class $\Lambda^{\rm all}$.

We now turn to the class $\Lambda^{\rm std}$. Without loss of generality we assume that $\lambda_2 < \lambda_1$ since otherwise $\mathbb{S}$ suffers from the curse of dimensionality also for the class $\Lambda^{\rm std}$. Then the choice of the element g for which $e_0(I_g) = e_0(S)$ in Theorem 1 is essentially unique and we take $g = g_d$ with $g_d = \lambda_{d,1}^{-1/2} S_d \eta_{d,1}$. We have $S_d^* g_d = \lambda_{d,1}^{1/2} \eta_{d,1}$ which is a tensor product since

$$S_d^* g_d = (\lambda_1^{1/2} \eta_1) \otimes \cdots \otimes (\lambda_1^{1/2} \eta_1) \quad (d \text{ times}).$$

Let

$$I_1 f = \left\langle f, \lambda_1^{1/2} \eta_1 \right\rangle_{F_1} \qquad \text{for all} \quad f \in F_1.$$

Then

$$I_d = I_1 \otimes \cdots \otimes I_1 \quad (d \text{ times})$$

is a linear tensor product functional. We have $I_d = I_{g_d}$ and $e_0(I_d) = e_0(S_d)$. Let

$$\mathbb{I} = \{I_d\}_{d=1}^{\infty}.$$

Theorem 1 yields that

$$n(\varepsilon, I_d, \Lambda^{\rm std}) \le n(\varepsilon, S_d, \Lambda^{\rm std}) \qquad \text{for all} \quad \varepsilon \in (0,1), \ d = 1, 2, \ldots .$$

This implies the following corollary

Corollary 4.
If one of the tractability notions does not hold for $\mathbb{I}$ then it also does not hold for $\mathbb{S}$. □

We now illustrate Corollary 4 for two examples for which $\mathbb{I}$ is multivariate integration and for which it is known that multivariate integration suffers from the curse of dimensionality.

Example 5. Korobov Space
As in [2, 9], let F_1 be a Korobov space whose reproducing kernel is

$$K_1(x,y) = 1 + 2\beta \sum_{k=1}^{\infty} \frac{\cos(2\pi(x-y))}{k^{2\alpha}} \qquad \text{for all} \quad x, y \in [0,1]$$

for some $\beta \in (0,1]$ and $\alpha > \frac{1}{2}$. This corresponds to the norm

$$\|f\|_{F_1}^2 = |\widehat{f}(0)|^2 + \beta^{-1} \sum_{h \in \mathbb{Z} \setminus \{0\}} |h|^{2\alpha} \, |\widehat{f}(h)|^2,$$

for Fourier coefficients $\widehat{f}(h)$ of f. We take $G_1 = L_2([0,1])$ and $S_1 f = f$, hence we consider the approximation problem $\mathbb{APPP}$.

In this case we know that

$$\lambda_1 = 1, \quad \lambda_2 = \beta \quad \text{and} \quad \eta_1 \equiv 1.$$

Hence, for $\beta = 1$ multivariate approximation suffers from the curse of dimensionality for Λ^{all} (and of course for Λ^{std}), and for $\beta < 1$, multivariate approximation is QPT with the exponent

$$t^* = \max\left(\frac{1}{\alpha}, \frac{2}{\ln\, \beta^{-1}}\right).$$

Note that $g_1 = \lambda_1^{-1/2} S_1 \eta_1 \equiv 1$. Therefore $g_d \equiv 1$ and

$$I_d(f) = \int_{[0,1]^d} f(x)\,\mathrm{d}x \quad \text{for all} \quad f \in F_d$$

is multivariate integration. From Theorem 16.16 on p. 457 in [7] which is based on [3, 5], we know that multivariate integration suffers from the curse of dimensionality. So does multivariate approximation due to Corollary 4. □

Example 6. Sobolev Space

We now take the Sobolev space F_1 of absolutely continuous functions on $[0,1]$ whose first derivatives are square integrable with the inner product

$$\langle f, u\rangle_{F_1} = \int_0^1 f(x)\,u(x)\,\mathrm{d}x \;+\; \int_0^1 f'(x)\,u'(x)\,\mathrm{d}x.$$

This space has the intriguing reproducing kernel

$$K_1(x,t) = \frac{1}{\sinh(1)}\,\cosh(1-\max(x,t))\,\cosh(\min(x,t)) \quad \text{for all} \quad x,t \in [0,1],$$

see [1]. We consider the L_2 approximation problem, as in Example 4. Hence we have $S_1 f = f$ and $G_1 = L_2([0,1])$.

In this case we have for $d = 1$ that $\lambda_1 = 1$ is of multiplicity 1, and $\eta_1 \equiv 1$. The second largest eigenvalue satisfies the condition

$$\lambda_2 = \max_{f \in F_1,\, \int_0^1 f(x)\,\mathrm{dx}=0} \frac{\int_0^1 f^2(x)\,\mathrm{d}x}{\int_0^1 f^2(x)\,\mathrm{d}x + \int_0^1 [f'(x)]^2\,\mathrm{d}x} = \frac{1}{1+\mu_2},$$

where

$$\mu_2 = \min_{f \in F_1,\, \int_0^1 f(x)\,\mathrm{d}x = 0} \frac{\int_0^1 [f'(x)]^2\,\mathrm{d}x}{\int_0^1 f^2(x)\,\mathrm{d}x}.$$

It is known, see e.g. [10], that $\mu_2 = \pi^2$. Hence, have

$$\lambda_2 = \frac{1}{1+\pi^2} = 0.091999668\ldots.$$

It is well known that $\lambda_j = \Theta(j^{-2})$ so that $\mathrm{decay}_\lambda = 2$. This implies that multivariate approximation for tensor products $F_d = F_1^{\otimes d}$ and $G_d = L_2([0,1]^d)$ is QPT for Λ^{all} with the exponent

$$t^* = 1.$$

Due to the form of (λ_1, η_1), the linear functional I_d corresponds to multivariate integration. It is known that multivariate integration suffers from the curse, see [11] which is also reported in [7] pp. 605-606. Hence, multivariate approximation also suffers the curse of dimensionality for the class Λ^{std} due to Corollary 4. □

Tractability of tensor product functionals $\mathbb{I}$ was thoroughly studied in [5], see also Chapters 11 and 12 of [7]. In particular, for many spaces F_1 the problem $\mathbb{I}$ suffers from the curse of dimensionality for the class Λ^{std}. This holds if the reproducing kernel K_1 of F_1 has a decomposable part and the univariate function η_1 has non-zero components with respect to the decomposable part. If this is the case then $\mathbb{S}$ also suffers from the curse of dimensionality for the class Λ^{std} although we may have QPT for the class Λ^{all}. We will mention more specific results in the next section.

5 Sobolev Space

We now consider tensor product problems $\mathbb{S}$ defined as in the previous section for the space F_1 taken as a Sobolev space of univariate real functions defined over $[0,1]$. More precisely, let F_1 be the space of absolutely continuous functions defined over $[0,1]$ and whose first derivatives belong to $L_2([0,1])$. The space F_1 has the reproducing kernel

$$K_1(x,y) = 1 + \min(x,y) \quad \text{for all} \quad x, y \in [0,1], \tag{4}$$

and the inner product for $f, h \in F_1$ is

$$\langle f, h \rangle_{F_1} = f(0)\,h(0) + \int_0^1 f'(x)\,h'(x)\,\mathrm{d}x.$$

For the tensor product space $F_d = F_1^{\otimes d}$ of d copies of F_1, the inner product for $f, h \in F_d$ is now of the form

$$\langle f, h \rangle_{F_d} = f(0)\, h(0) + \sum_{\emptyset \neq \mathfrak{u} \subseteq \{1,2,\ldots,d\}} \int_{[0,1]^{|\mathfrak{u}|}} \frac{\partial^{|\mathfrak{u}|} f}{\partial x_{\mathfrak{u}}}(x_{\mathfrak{u}}, 0) \frac{\partial^{|\mathfrak{u}|} h}{\partial x_{\mathfrak{u}}}(x_{\mathfrak{u}}, 0)\, \mathrm{d}x_{\mathfrak{u}},$$

where $\partial x_{\mathfrak{u}} = \prod_{j \in \mathfrak{u}} \partial x_j$, $\mathrm{d}x_{\mathfrak{u}} = \prod_{j \in \mathfrak{u}} \mathrm{d}x_j$ and $(x_{\mathfrak{u}}, 0)$ is a d dimensional vector with components x_j for $j \in \mathfrak{u}$ and 0 otherwise.

It was proved in [7] pp.195-200, see also [8], that for any linear non-zero tensor product functional its information complexity (for Λ^{std}) is 1 or it is exponentially large in d. Furthermore, the information complexity is 1 only for trivial cases when the linear tensor product functional is of the form

$$a^d f(t, t, \ldots, t) = \langle f, h_d \rangle_{F_d} \quad \text{with} \quad h_d(x) = \prod_{j=1}^{d} a\,(1 + \min(x_j, t))$$

for some non-zero real a and for some $t \in [0, 1]$. Applying this results for $\mathbb{I}$ we see that as long as

$$\eta_1 \neq a(1 + \min(\cdot, t)) \quad \text{for all } a \in \mathbb{R} \text{ and } t \in [0, 1] \tag{5}$$

then $\mathbb{I}$ as well as $\mathbb{S}$ suffer from the curse of dimensionality for the class Λ^{std}. We summarize the results from the last two sections in the following theorem.

Theorem 7.

Consider a linear non-zero tensor product problem $\mathbb{S}$, as defined in this section.

- Let $\lambda_2 = \lambda_1$. Then $\mathbb{S}$ suffers from the curse of dimensionality for Λ^{all} (and Λ^{std}).
- Let $\lambda_2 < \lambda_1$. Then $\mathbb{S}$ is QPT for Λ^{all} iff (3) holds.
- Let $0 < \lambda_2 < \lambda_1$. Then $\mathbb{S}$ is not PT for the class Λ^{all}.
- Let $\lambda_2 < \lambda_1$. If (5) holds then $\mathbb{S}$ suffers from the curse of dimensionality for Λ^{std}.

□

In general, the assumption (5) used in the last part of Theorem 7 is needed. Indeed, if (5) does not hold then we may have S_d as a linear tensor product functional of the form $S_d f = a^d f(t, t, \ldots, t)$ with a nonzero real a and $t \in [0, 1]$. Then $\mathbb{S}$ is trivial since $n(\varepsilon, S_d, \Lambda^{\mathrm{std}}) = 1$ for all $\varepsilon \in [0, 1)$ and all d.

We now verify the assumptions (3) and (5) for multivariate approximation $\mathbb{APP}$.

The eigenpairs (λ_j, η_j) were found in [13], see also [9] pp. 409-411. We have $\lambda_j = \alpha_j^{-2}$, where $\alpha_j \in ((j-1)\pi, j\pi)$ is the unique solution of the nonlinear equation

$$\cot\, x = x \qquad \text{for} \;\; x \in ((j-1)\pi, j\pi),$$

and

$$\eta_j(x) = \beta_j \, \cos(\alpha_j x - \alpha_j) \qquad \text{for all} \;\; x \in [0,1],$$

where

$$\beta_j = \left(\cos^2(\alpha_j) + \tfrac{\alpha}{2} \, \left(\alpha_j - \tfrac{1}{2}\sin(2\alpha_j)\right)\right)^{-1/2}.$$

For $j = 1$ and $j = 2$ the numerical computation yields

$$\begin{aligned} \lambda_1 &= 1.35103388\ldots, \\ \lambda_2 &= 0.08521617\ldots. \end{aligned}$$

Clearly, $\alpha_j = \pi\, j\, (1 + o(1))$ as j tends to infinity. This shows that

$$\lambda_j = \frac{1 + o(1)}{\pi^2\, j^2} \qquad \text{as} \;\; j \to \infty.$$

Therefore $\text{decay}_\lambda = 2$. Hence (3) holds and $\mathbb{APP}$ is QPT for the class Λ^{all}.

The assumption (5) holds if for all $a \in \mathbb{R}$ and $t \in [0,1]$ we can find $x \in [0,1]$ such that

$$\cos(\alpha_1 x - \alpha_1) \neq a(1 + \min(x,t)).$$

For $t = 0$, the right hand side is constant, whereas the left hand side varies. For $t > 0$, the right hand side is constant over $[0,t]$, whereas the left hand side varies. Therefore (5) also holds. We summarize this in the following corollary.

Corollary 8.

Consider the multivariate approximation problem $\mathbb{APP}$ for the Sobolev spaces studied in this section. Then

- $\mathbb{APP}$ is QPT but not PT for the class Λ^{all} with the exponent of QPT

$$t^* = 1.$$

- $\mathbb{APP}$ suffers from the curse of dimensionality for the class Λ^{std}.

□

We add in passing that a similar analysis can be done if the reproducing kernel (4) is replaced by

$$K_{1,a}(x,t) = 1 + \tfrac{1}{2}\left(|x-a| + |t-a| - |x-t|\right) \quad \text{for all } \; x,t \in [0,1],$$

for any $a \in [0,1]$ or by

$$K_1(x,t) = \min(x,t) \quad \text{for all } \; x,t \in [0,1].$$

For these variants of the Sobolev spaces, Corollary 8 is valid.

Remark 9. We conclude this paper with a comment on the rates of convergence and tractability notions for $\Lambda^{\rm all}$ and $\Lambda^{\rm std}$. In [4], the L_2 approximation problem was studied. It was shown that there are classes F for which the best rate of convergence of algorithms using n appropriately chosen linear functionals is $n^{-1/2}$ whereas for n function values the best rate can be arbitrarily bad. If the best rate for $\Lambda^{\rm all}$ is faster than $n^{-1/2}$ than we still do not know whether the rates for $\Lambda^{\rm std}$ and $\Lambda^{\rm all}$ always coincide.

For the examples in this section the rates for $\Lambda^{\rm std}$ and $\Lambda^{\rm all}$ are basically (up to log terms) n^{-1} but tractability properties for $\Lambda^{\rm all}$ and $\Lambda^{\rm std}$ are quite different. Hence, even if the rates are the same, tractability properties can be quite different for $\Lambda^{\rm all}$ and $\Lambda^{\rm std}$.

Acknowledgement. We thank Mario Ullrich who computed some numbers for us. We also thank Greg Wasilkowski, Markus Weimar and two referees for valuable comments.

References

[1] A. Berlinet and C. Thomas-Agnan, *Reproducing Kernel Hilbert Spaces in Probability and Statistics*, Kluwer Academic Publishers, Boston, 2004.

[2] M. Gnewuch and H. Woźniakowski, Quasi-polynomial tractability, *J. Complexity* 27, 312-330, 2011.

[3] F. J. Hickernell and H. Woźniakowski, Tractability of multivariate integration for periodic functions, *J. Complexity* 20, 660-682, 2001.

[4] A. Hinrichs, E. Novak and J. Vybiral, Linear information versus function evaluations for L_2-approximation, *J. Approx. Th.* 153, 97-107, 2008.

[5] E. Novak and H. Woźniakowski, Intractability results for integration and discrepancy, *J. Complexity* 17, 388-441, 2001.

[6] E. Novak and H. Woźniakowski, *Tractability of Multivariate Problems*, Volume I: Linear Information, European Math. Soc. Publ. House, Zürich, 2008.

[7] E. Novak and H. Woźniakowski, *Tractability of Multivariate Problems*, Volume II: Standard Information for Functionals, European Math. Soc. Publ. House, Zürich, 2010.

[8] E. Novak and H. Woźniakowski, Tractability of approximating multivariate linear functionals. In: The Steve Smale Festschrift, *J. Fixed Point Th. Appl.*, 7, 313-324, 2010.

[9] E. Novak and H. Woźniakowski, *Tractability of Multivariate Problems*, Volume III: Standard Information for Operators, European Math. Soc. Publ. House, Zürich, 2012.

[10] L. E. Payne and H. F. Weinberger, An optimal Poincaré inequality for convex domains. *Arch. Rational Mech. Anal.*, 5, 286-292, 1960.

[11] I. H. Sloan and H. Woźniakowski, Tractability of integration in non-periodic and periodic weighted tensor product Hilbert spaces, *J. Complexity* 18, 479-499, 2002.

[12] J. F. Traub, G. W. Wasilkowski and H. Woźniakowski, *Information-Based Complexity*, Academic Press, 1988.

[13] G. W. Wasilkowski and H. Woźniakowski, Weighted tensor-product algorithms for linear multivariate problems, *J. Complexity* 15, 402–447, 1999.